

Social Boundaries of Appropriate Speech in HCI: A Politeness Perspective

Leigh Clark
University College Dublin
School of Information & Communication Studies, Dublin, Ireland
leigh.clark@ucd.ie

This position paper presents the view that there may be social boundaries in appropriate speech in HCI. While previous research has examined that humanlike voices may not always be appropriate for computers to use, the same may hold true for linguistic concepts that have noticeable interpersonal and social functions. This paper examines the social functions of linguistic politeness to approach the issue of appropriate language use in spoken HCI, and the relationship between voices and language in this interaction space. Several studies exploring politeness in spoken HCI are discussed, with the view that it and other social talk may have different limitations in HCI, and that politeness itself may have to be subsequently reconsidered.

Politeness, appropriateness, identity, speech, voice user interface

1. INTRODUCTION

Speech contains a wealth of interactional complexities that build and maintain the way people see each other in terms of power, identity, and personality (Cameron, 2001; Coulthard, 2013; Goffman, 2005). What we say and how we say it can impact on how we are perceived by others. This can also be seen in spoken human-computer interaction (HCI). However, despite the increasing prevalence of computer speech, there is still debate as to how we may want computers to speak to us. Moore (2017a), for example, argues that the frequency of humanlike versus robotic voices in computer speech is not necessarily appropriate. This is due to people expecting that systems have more advanced capabilities than they do in reality. This can create a mismatch between users' mental models of what a system is capable of and the reality of its limitations (Cowan et al., 2017).

As well as with voice, there remain questions as to what could be considered appropriate language in computer speech. There are voices that are considered humanlike and machinelike or robotic. While all language is arguably humanlike, some concepts may be more appropriate for computers to use than others or be considered 'machinelike'. This paper uses the concept of linguistic politeness as a lens to examine the possible social boundaries in computer language use during spoken interaction, and the relationship between language and voice in perceptions of computer speech. Subsequent challenges surrounding concepts of

gender, vocal qualities, appropriateness, trust, and interaction contexts are discussed alongside reconsidering politeness in HCI contexts.

2. POLITENESS IN HCI

While there is wider debate around what constitutes politeness (Locher & Watts, 2005) and its limitations outside of Anglo-centric countries, this paper looks at examples of positive and particularly negative politeness strategies as described by Brown & Levinson (1987). These strategies are often linked to the concept of *face* – the public self-image we present to others during social interaction (e.g. Goffman, 2005). In face theory, it is often in the best interest of interlocutors to *save face* during interaction. This means avoiding damage to one's own or another's self-image. This is achieved through a dynamic understanding of what is acceptable and unacceptable communication at any given moment. This understanding can be impacted by various societal and cultural norms, which are themselves subject to change. Brown & Levinson (1987) described *positive face* – the desire to be liked and approved – and *negative face* – the desire to not be imposed upon by others. Accordingly, they describe positive and negative politeness strategies that speakers use to conduct face-saving in interaction. Positive politeness strategies include approval, attention to a listener's wants and desires, and indication of group membership between interlocutors. Negative politeness strategies often refer to minimising the imposition

made on others, for example by using indirect instructions or requests to others.

In HCI research, the use of politeness strategies has had mixed responses from participants. Wang et al. (2008) demonstrated that polite pedagogical agents used in the classroom may lead to positive improvements in learning outcomes. More positive effects were observed in a study where participants observed interactions between people and robot instructors. Torrey, Fussell & Kiesler (2013) observed improvements in perceptions of likeability and considerateness, as well as reductions in how controlling a robot was perceived to be, when discourse markers and hedges were used.

This experiment design was expanded upon, where the interaction distance between robot and participant was further explored to compare direct and indirect interactions, and observations of interactions with polite robots (Strait, Canning, & Scheutz, 2014). In this study, the authors found the same improvements in perceptions were only observed in third-person interactions, similar to the observations described by Torrey et al. (2013). Direct interactions with polite robots did not show the same effects.

Positive improvements have also been observed in the wild. While politeness as a concept may not have been the direct focus of their research, a large supermarket chain used politeness strategies when updating the language of a self-service checkout in (Clark, 2016). For example, the phrase “*Approval needed*” was replaced with “*We just need to approve this*”. The use of the hedge *just* can be seen to minimise the perception of what constitutes the approval, and the potential time it may take to accomplish. While perhaps not entirely a result of employing politeness strategies, improvements in people’s user experience with the checkout systems were noted.

Although mixed responses were observed when using politeness in HCI, factors contributing to these responses were not always made clear. Two further studies provided more insight into perceptions of politeness in HCI, in which participants interactions with computer instructors were examined during model building tasks. Negative politeness strategies were used in the form of vague language (Channell, 1994; Cutting, 2007). The first study saw polite instructions being perceived as inappropriate in synthesised speech instructions compared to non-polite instructions, which were perceived as relatively normal and expected (Clark, Bachour, Ofemile, Adolphs, & Rodden, 2014). The second study showed a marked improvement in how polite instructions were received when used in human-recorded instructions for a computer interface, compared with synthesised speech (Clark, Ofemile, Adolphs, & Rodden, 2016). The polite recorded voice was

perceived as more appropriate in using polite instructions than two synthesised voices. However, there were still noticeable limitations in its appropriateness. Even with the human recordings, participants commented on how the interface was still “just a machine” and wasn’t capable of using politeness in the same way as other people.

3. APPROPRIATENESS AND SOCIAL BOUNDARIES

These findings raised the question of possible boundaries with what is considered appropriate computer speech. While participants in the studies could not always explicitly identify the language that was causing the negative reactions, they could identify a disparity between the language used and the interface that was delivering it. The expectations of appropriate and acceptable language use of a more robotic sounding voice appeared to be more limited than a humanlike voice. This is similar to the gap between reality and expectations observed when comparing robotic-sounding and human voices (Moore, 2017a), and may represent somewhat of a verbal uncanny valley resulting in a sense of unease.

However, given the mixed reactions towards even the humanlike voice using politeness, there may be further limitations to appropriate spoken communication in HCI. As Moore (2017b) discusses, there may be limits as to what linguistic interactions can take place between humans and machines as unequal partners. While a human voice may help blur boundaries between expected human speech and expected computer speech, and act as a bridge between identities of sorts, there was still somewhat of a clash of users’ mental models in the expectation of linguistic capabilities. What is expected and possible of a human may not automatically transfer to a computer.

3.1. Differences in social rules

Considering the link between the use of politeness strategies and face management, we may have to consider that the idea of face in HCI is not the same as in human interaction. This may also be true for social rules of communication. While people have been observed to attribute politeness to computers (Nass, Steuer, & Tauber, 1994) and formulate polite responses towards them (Large, Clark, Quandt, Burnett, & Skrypchuk, 2017), responses to politeness particularly in instruction-based contexts are varied. This is to be expected to an extent, as responses to language are also diverse in human communication. However, it may be the case that face threats do not carry the same weight in HCI as they do in human communication. Although a mutual acknowledgement of self-images and potential face threats is said to exist between people, the same might not be true between people and computers. Users still have an

autonomy that can be imposed upon, and there are stakeholders involved (often commercial) for the computer, but likely not to the same extent as with other people.

4. RETHINKING POLITENESS IN HCI

Politeness in human communication is debated. Some see it as too focused on the polite end of a much larger polite-to-impolite spectrum and consider impolite language just as important in interaction. This is described, for example, in the concept of *relational work* – all aspects of interaction that contribute towards building and maintaining interpersonal and social relationships (Locher & Watts, 2005). Politeness research in HCI has often focused on the use of negative politeness strategies. However, while researchers may benefit in expanding the concept to include all elements of the (im)polite spectrum, there are still considerations for interpersonal and rapport management through language in spoken HCI that differ from human interaction.

4.1. Third party involvement

Firstly, there are third party stakeholders at play, such as designers, that are involved in interactions to varying degrees. This can include choosing all system output explicitly, for example with the limited phrases used by self-checkouts or allowing for output to be determined from a bank of possible utterances. In human-human interactions, the explicit creation of linguistic output for others does exist (e.g. speech writers for politicians or business figureheads). However, this is less common than in HCI. The involvement of a third party during speech-based HCI may impact upon the evaluation of relational work and social talk like politeness strategies. The use of humour, for example, may be attributed to a system developer or company rather than the system itself. Devices may also be associated with companies and their respective reputations that are perceived by individuals may influence users' perceptions of computer language use.

4.2. Ownership and personalisation

Secondly, there is the aspect of ownership to consider. People can interact with systems like automated checkouts that are based within a private business. They can also interact with IPAs that are part of their personal property, be it on a smartphone or a home-based intelligent assistant. There may be a personal sense of attachment to a smartphone, for example, or a distrust of devices placed within corporate spaces. This also ties in with the concept of personalisation – being able to customise a device to one's own preferences – the availability of which may depend on whether or not the user owns the device in question. Being able to alter characteristics such as the gender of a

computer's voice may alter the way in which language is perceived by users.

5. FUTURE WORK AND CHALLENGES

In reconsidering politeness and other social talk in HCI, and the factors that affect its use and perception during interaction, a number of challenges arise for future research. These are briefly outlined below.

5.1. Gender and voice quality

One feature often absent from research discussed in this paper is a comparison between politeness in male and female synthesised speech. Some research suggests politeness is more common amongst female speakers than male speakers (Lakoff, 2004), though this stereotype has been challenged over the years (Mills, 2005). Nevertheless, the ability to personalise gender of voices in HCI with some devices presents an interesting avenue of exploration – expectations and perceptions of politeness may be affected by both a system's perceived gender and that of the user. This may be confounded by the availability of choice. The relationship between language and a wealth of other vocal qualities also presents itself as further areas of research, both within and outside of device personalisation. Theories including the *similarity-attraction effect* (Nass & Lee, 2000) could be tested alongside language use, focusing on qualities such as accents and dialects. Similarities between users and systems may affect the perception of concepts like politeness in speech-based HCI.

5.2. Measuring appropriateness and trust

The measurement of appropriateness of politeness and social talk in HCI remains another challenge. Moore (2017a) provides some examples of how appropriateness may be measured with computer voices. However, language use is arguably a more complex affair. Focusing on specific linguistic concepts (e.g. politeness, face, vague language) would likely prove an easier way of analysing appropriateness, as opposed to the essentially infinite utterances available. The link between appropriateness of language and trust would also prove useful, particularly in contexts that require systems to interrupt their users and direct the flow of conversation. In such instances, trust of would be of great importance. Understanding what linguistic concepts would influence this would be highly valuable. Further experiments can test the effect of politeness and other social talk during system-led interruptions, and how this may impact user perceptions of appropriateness and trust of information and systems in these contexts.

5.3. Moving beyond instructions and ‘one-shot’ interactions

The appropriateness of politeness may lie in the context of use. Indeed, exploring a wider array of contexts is essential given the unprecedented level of available spoken interactions with computers. A number of the studies discussed in this paper investigate instruction or advice-giving contexts, usually with static interfaces or robots. The effects of social talk in other contexts, including mobile contexts and more collaborative environments where users and computers operate more at a peer level, would provide a richer understanding of how people like to be talked to by computers.

Furthermore, how we like to talk to computers ourselves should also be investigated with regards to social talk. Prior research has observed people using politeness strategies and vague language towards computers (Large et al., 2017), though it is less clear why this occurs. For instance, it may be the case of lexically aligning (i.e. matching the same language) with computers (Branigan, Pickering, Pearson, & McLean, 2010) but there may be further questions as to how beliefs about a system affect it occurring (Branigan, Pickering, Pearson, McLean, & Brown, 2011).

The research discussed in this paper is also relatively short-lived. While speech interfaces like IPAs are still relatively unused (Cowan et al., 2017) we nevertheless have a greater potential of using them across a longer period of time than in previous years. Longitudinal studies are few in this area of research, and interactions with other concepts including appropriateness and trust would be helpful for considering linguistic aspects of system design.

6. SUMMARY

This paper presents the view that there may be social boundaries to appropriate use of linguistic concepts such as politeness in spoken HCI. While negative politeness strategies, for example, have been perceived more positively when used by human versus synthesised voices, this picture is incomplete. The very nature of being a computer may limit its ability to appropriately and capably employ certain linguistic concepts that are inherently social. Using politeness as a lens, this paper reflects on the possible social boundaries surrounding appropriate computer speech. Considerations for rethinking politeness and social talk are presented, alongside challenges surrounding future work around gender, voice quality, appropriateness, trust, and interaction context. Future experiments are planned to assess these concepts. These will include developing a bottom-up understanding of what users expect from speech interfaces in terms of language and voice, and the how this is affected by different interaction

contexts. Furthermore, users’ interactions with a speech interface that employs politeness strategies will be assessed to investigate how synthetic and human speech affects people’s alignment to politeness strategies. These experiments will further the body of knowledge of politeness and social communication with speech interfaces, and further explore where boundaries of humanlike and machinelike communication may lie.

ACKNOWLEDGEMENTS

This research was funded by a New Horizons grant from the Irish Research Council entitled “The COG-SIS Project: Cognitive effects of Speech Interface Synthesis” (Grant R17339).

7. REFERENCES

- Branigan, H. P., Pickering, M. J., Pearson, J., & McLean, J. F. (2010). Linguistic alignment between people and computers. *Journal of Pragmatics*, 42(9), 2355–2368. <https://doi.org/10.1016/j.pragma.2009.12.012>.
- Branigan, H. P., Pickering, M. J., Pearson, J., McLean, J. F., & Brown, A. (2011). The role of beliefs in lexical alignment: Evidence from dialogs with humans and computers. *Cognition*, 121(1), 41–57. <https://doi.org/10.1016/j.cognition.2011.05.011>.
- Brown, P., & Levinson, S. C. (1987). *Politeness: Some Universals in Language Usage*. Cambridge University Press.
- Cameron, D. (2001). *Working with Spoken Discourse*. SAGE.
- Channell, J. (1994). *Vague Language*. Oxford University Press.
- Clark, L., Ofemile, A., Adolphs, S., & Rodden, T. (2016). A Multimodal Approach to Assessing User Experiences with Agent Helpers. *ACM Trans. Interact. Intell. Syst.*, 6(4), 29:1–29:31. <https://doi.org/10.1145/2983926>.
- Clark, Leigh M. H. (2016). *Exploring vague language use and voice variation in human-agent interaction* (unpublished Ph.D.). University of Nottingham. Retrieved from <http://eprints.nottingham.ac.uk/37863>.
- Clark, Leigh Michael Harry, Bachour, K., Ofemile, A., Adolphs, S., & Rodden, T. (2014). Potential of imprecision: exploring vague language in agent instructors (pp. 339–344). ACM Press. <https://doi.org/10.1145/2658861.2658895>.

- Coulthard, M. (2013). *Advances in Spoken Discourse Analysis*. Routledge.
- Cowan, B. R., Pantidi, N., Coyle, D., Morrissey, K., Clarke, P., Al-Shehri, S., ... Bandeira, N. (2017). 'What can i help you with?': infrequent users' experiences of intelligent personal assistants (pp. 1–12). ACM Press. <https://doi.org/10.1145/3098279.3098539>.
- Cutting, J. (Ed.). (2007). *Vague language explored*. Basingstoke [England]; New York: Palgrave Macmillan.
- Goffman, E. (2005). *Interaction Ritual: Essays in Face to Face Behavior*. Routledge.
- Lakoff, R. T. (2004). *Language and Woman's Place: Text and Commentaries*. Oxford University Press.
- Large, D. R., Clark, L., Quandt, A., Burnett, G., & Skrypchuk, L. (2017). Steering the conversation: A linguistic exploration of natural language interactions with a digital assistant during simulated driving. *Applied Ergonomics*, 63, 53–61. <https://doi.org/10.1016/j.apergo.2017.04.003>.
- Locher, M. A., & Watts, R. J. (2005). Politeness Theory and Relational Work. *Journal of Politeness Research. Language, Behaviour, Culture*, 1(1). <https://doi.org/10.1515/jplr.2005.1.1.9>.
- Mills, S. (2005). Gender and impoliteness. *Journal of Politeness Research*, 1(2), 263–280. <https://doi.org/10.1515/jplr.2005.1.2.263>.
- Moore, R. K. (2017a). Appropriate Voices for Artefacts: Some Key Insights. 1st International Workshop on Vocal Interactivity in-and-between Humans, Animals and Robots.
- Moore, R. K. (2017b). Is Spoken Language All-or-Nothing? Implications for Future Speech-Based Human-Machine Interaction. In *Dialogues with Social Robots* (pp. 281–291). Springer, Singapore. https://doi.org/10.1007/978-981-10-2585-3_22.
- Nass, C., & Lee, K. M. (2000). Does computer-generated speech manifest personality? an experimental test of similarity-attraction (pp. 329–336). ACM Press. <https://doi.org/10.1145/332040.332452>.
- Nass, C., Steuer, J., & Tauber, E. R. (1994). Computers are social actors (p. 204). ACM Press. <https://doi.org/10.1145/259963.260288>.
- Strait, M., Canning, C., & Scheutz, M. (2014). Let me tell you! investigating the effects of robot communication strategies in advice-giving situations based on robot appearance, interaction modality and distance (pp. 479–486). ACM Press. <https://doi.org/10.1145/2559636.2559670>.
- Torrey, C., Fussell, S. R., & Kiesler, S. (2013). How a robot should give advice (pp. 275–282). IEEE. <https://doi.org/10.1109/HRI.2013.6483599>.
- Wang, N., Johnson, W. L., Mayer, R. E., Rizzo, P., Shaw, E., & Collins, H. (2008). The politeness effect: Pedagogical agents and learning outcomes. *International Journal of Human-Computer Studies*, 66(2), 98–112. <https://doi.org/10.1016/j.ijhcs.2007.09.003>.